

General Options - One of the most perplexing dilemmas confronting configuration planning today is the evolving technology and economy of memories. We have shown our planned configuration incorporating strictly DEC equipment including memory for simplicity. We are well aware, however, that DEC memory is not currently economically competitive and as indicated earlier in the small machine support proposal, the cost of large electronic memories in general is precarious and likely to drop significantly in the next few years. Whereas we must obtain sufficient memory immediately with our machine to allow initial developments, we will carefully consider a variety of manufacturers, incremental expansion, and leasing so as to optimize our posture for taking advantage of these new memory developments as they become available. A corollary problem exists in the peripheral equipment field such as disk drives.

In addition to these component problems several options exist in the design of terminal interfaces and synchronous data interfaces with the PDP-10. These trade-offs involve the relative cost and technical desirability of using a minicomputer interface which could require system software modification as opposed to standard DEC interface systems.


In general, the specification of a particular approach here does not preclude the lease or purchase of an alternate as long as system performance is optimized, the cost is attractive, and the agencies and regulations permit. In fact every effort will be made to search the market for the most effective and economical alternates, consistent with design goals.

6. Implementation Plan


The implementation of the proposed dual machine facility must proceed so as to minimize the impact on on-going service computing. The phased conversion effort is shown schematically in Figure 5. Under this grant, the first PDP-10 machine will be procured and the PL/ACME software converted to run utilizing the PDP-10 time-share monitor functions and relocation features. When this system is checked out, the 360/50 will be removed with the service computing function wholly transferred to a second PDP-10 installed to take its place.

The facility accommodations for the two machines can be made within a straightforward modification of the current 360/50 facility. This modification is consistent with currently approved building plans for a corresponding arcade in the Medical Center; no formal commitment to these modifications has yet been requested from the University.

I. HARDWARE TIME LINE



II. SOFTWARE TIME LINE



IMPLEMENTATION PHASING PLAN

FIGURE E-5

*Assumes the selection of a PDP-10 as the service machine.

F. RESOURCE PERSONNEL

The senior personnel associated with this grant's core research activity will include Dr. Joshua Lederberg, Principal Investigator, Dr. Edward Feigenbaum Associate Investigator, and the technical staff with considerable experience in the development of computing techniques.

Among the technical staff currently employed in biomedical computing is Thomas Rindfleisch who supervised image processing development in an applications group in the Jet Propulsion Laboratory prior to coming to Stanford. He played a central role in developing the image processing for the Mariner Mars 1969 Space Photographic missions, which has also been extended to the realtime analysis of Mars photographs in the current (1971-72) Mariner program. These studies have also led to on-going advances at JPL in image enhancement of X-ray images, light microscope images, and other biomedical applications. At Stanford he has been engaged in the design of extended realtime systems for the DENDRAL Project and in planning activities for Medical Center computing in general. His nine years of experience in digital image processing at JPL provide a solid base of capability with which to undertake the imaging portions of the collaborator's projects.

Another member of the technical staff is Gio Wiederhold who led the design and implementation of the ACME time-sharing system. After directing the ACME facility for its initial five years, he participated in the design of a new computer architecture and consulted regularly with biomedical computer users. In the recent past he has assisted a number of new realtime users on the ACME system with their data acquisition problems.

Another senior member of the technical staff is Lee Hundley, who implemented much of the realtime data collection system under ACME. In addition, Lee Hundley has written the assemblers for PDP-8's and PDP-11's which are currently available to ACME users and has supervised the systems development group during the past three years.

Dr. Walter Reynolds has been working in the Instrumentation Research Laboratory for several years and has authored a number of publications on realtime connections of laboratory instruments to computers.

Other members of the ACME computing facility have significant experience in communications, file handling, compiler improvement, graphics, small machine assemblers and simulators, and related fields. Important contributions to the proposed goals can be made by a number of well-experienced staff members.

One member of the systems programming staff will **serve** as a point of contact on each collaborator's project. We would expect these programmers' loyalties to rest with the collaborator's objectives. It is felt that such an assignment would help to produce good communication and cooperation.

The collaborators who will perform the research for this proposed grant are a group of sophisticated, experienced computer users. Both the medical staff and the programming staff have achieved notable success in the past. The programming staff in the Division of Cardiology, led by William J. Sanders (formerly of the ACME staff), has achieved notable success in realtime monitoring of cardiac catheterization procedures and video image analysis. Chris Constantinou in Urology has spent the past three years studying the movement of the ureter using animal subjects connected in realtime to the ACME system. The Instrumentation Research Laboratory in the Department of Genetics has made major contributions to the DENDRAL and cell separation projects. Included in the IRL

engineering and programming staff are Dr. Walter Reynolds and Mark Stefik. Additional examples of accomplishments by current staff are briefly described in the collaborators' research statements and the lists of publications attached to the biographical sketches.

Another example of highly significant computer related research at Stanford is DENDRAL. This project has achieved acclaim as an artificial intelligence project performing automated interpretation of mass spectra. A number of publications authored by this group are cited in the collaborative project description and the individual biographical sketches. This research team adds considerable depth to the measure of technical talent affiliated with the proposed resource.

Over these past six years, ACME has had a number of outstanding successes. The initial implementation of a time-shared interactive complex in eighteen months was a feat of some magnitude. The accomplishment of a medium data rate, realtime data acquisition and control system in the framework of a time-shared system is impressive. ACME's major accomplishment, the education of the user community, can best be appreciated by noting that over 1000 people have attended ACME courses in the past five years. These accomplishments are attributable to the excellence, dedication and motivation of the ACME staff.

G. RESOURCE ACTIVITIES

1. Services to be Provided

This proposal calls for the creations of a research resource called SUMEX. SUMEX will offer services of four types:

- a. Central computing hardware (primarily in the form of a DEC PDP-10 system with KI-10 processor, 192K words of main memory and 4 satellite PDP-11's connected via a PDP10/11 interface. See figure E-3.) for use in computer science related research in biomedicine and development of new hardware interfaces to users' peripheral minicomputers.
- b. Development of software support under a core research program and software support for collaborative research projects.
- c. The resources will provide backup computing facilities for the Stanford University Medical Center Computer Facility.
- d. The resource will include a group of people with strong technical skill to support computer-science type research in the biomedical community. Some additional personnel resources which can be shared among several collaborators are planned; the applied mathematician in the budget is one such position.

The SUMEX resource will support research on advanced computer applications. Specific examples of support anticipated are listed in the core research and collaborative projects described below.

2. Research

a. Core Research

(1) Satellite Machine Support

(a) Problem Statement

- (1) Limitations of stand-alone small computers. The primary limitations of existing stand-alone systems are a) the economics of primary and secondary storage and b) the lack of highly sophisticated software. Secondary storage costs for small stand-alone systems tend to run twice the corresponding costs on large systems. Primary storage (core) costs approximately 2¢ per bit for Ampex bulk core memory (i.e., for the 360/50) capable of running at approximately 2.4 microseconds. DEC supplied 1.8 microsecond memory for the PDP-11 costs approximately \$3,000 for 4K words or 4.5¢ per bit. The second major limitation of stand-alone small systems is the limited support for sophisticated programming languages and interactive capabilities. This implies that large amounts of effort are spent for software development even in rather simple applications. Some additional limitations are as follows:

- (a) Because the hardware investment is low, users seem to have difficulty justifying significant programming effort. Only users with a large number of similar machines can justify a substantial software systems effort.
- (b) The manufacturers tend to scale their system software to the smaller end of the line.
- (c) Owners of small machines lack easy access to one another's data base because their planning has assumed dedicated use of the hardware.
- (d) Finally, the software systems do not provide the ability to concurrently utilize a variety of system capabilities. For example, one cannot be performing realtime data acquisition and editing programs simultaneously.

(2) Attributes of large host systems.

- (a) Needs not met by satellite systems. A large host computer can provide an array of services not available to stand-alone small machines in today's technology. Among such services are:
 - (1) access to large volumes of source programs,
 - (2) ability to edit them efficiently,
 - (3) ability to concurrently use multiple system capabilities,
 - (4) availability of highly sophisticated programming languages and interactive capabilities,

- (5) availability of broader line of peripherals,
 - (6) shared access to common databases, and
 - (7) more extensive software support from the hardware manufacturer.
- (b) Special attributes of larger systems. Larger systems such as the PDP-10 can directly address large quantities of memory (the PDP-10 can address directory up to 256K words). The large system has more input/output ports which can operate concurrently. Large machines provide a standard capability for handling scientific calculations (floating point hardware, etc.). Hardware and software systems are currently available for measuring performance on large systems; this cannot be said for most small systems.
- (c) Impact on users. The objective of this part of the proposal is to provide to the small stand-alone user a link to a larger system. The link will make available to the satellite system services which can be obtained today on larger systems but cannot be obtained today on smaller systems. Thus, the satellite user will have available more sophisticated programming languages, access to cheaper primary and secondary storage, shared database, shared communications systems, backup support for large computational problems, and an overall significant degree of labor saving. In today's computer environment, most users can expect to spend two or more times their hardware investment for software development. The trend of this curve is to increase software costs relative to hardware (some say geometrically).
- Small machine users have adopted patterns of thought which were appropriate to dedicated stand-alone systems. Part of the mission of this grant will be to demonstrate the feasibility of broader services which could be available in satellite mode and to expand their decision-making analysis to include more flexible and sophisticated computing services. If successful, the program should also enable the medical user community to save a considerable cost in hardware and software. Most satellites would no longer require an extensive collection of peripheral hardware.
- (3) Anticipated direction of some technological innovations. Given the rapidly expanding market for computer-based applications and cost-effective hardware systems, major development effort is being expanded in the industry for the development of new techniques. Intelligent planning requires cognizance of the most likely near-term advances in relevant technology, involving the following:
- (a) Primary storage. Very low cost primary storage will become available quite soon. The costs associated with integrated circuit memory can be expected to reach less than 1/10 the cost of the current MOS systems. The primary reason for this significant drop in cost is the removal of the high labor content of current core technologies. Currently, MOS costs are close to those of core memories, despite the fact that it is a very recent development and still heavily burdened with development costs.

- (b) Secondary storage. Secondary storage systems are being developed which will have faster access times and transfer rates and higher density on the storage media. Current secondary storage technology involves physically rotating memories with inherent mechanical complexity and high costs of achieving the tolerances required and of maintaining the the system once installed. Major breakthroughs in the current technology bottleneck appear possible through development of magnetic bubble memories and charge coupled devices. In terms of information made public to date, Bell Telephone Laboratories seems to be spending the largest effort in both of these new non-mechanical approaches to large data storage devices. Corollary efforts on the parts of other manufacturers have a very low profile at this time.
- (c) Low cost general purpose processors. The recent development of large scale integrated circuitry makes possible today the building of general purpose processors at very low costs. For example, INTEL Corporation now has an 8 bit central processor (designed with MOS technology) available on a single chip. The cost of such a component is expected to drop below \$50 within the next year. It is of greater sophistication than a PDP-8, although 10 times slower. We anticipate that such developments will lead to new small computer organizations incorporating large numbers of such central processing units dedicated to a variety of applications. Of course this is but one element in a system, the balance of which includes significant cost components. The impact of this change cannot be postponement of solutions needed now. We can cite two recent examples of the development of special purpose processors which support applications on a general purpose computer. The Berkeley Computer Corporation, BCC-1, incorporated a number of micro-processors which were assigned tasks such as:
- (1) file management,
 - (2) drum scheduling for a large swapping drum,
 - (3) terminal input/output and control of remote communication links.

In order to process large volumes of experimental data in realtime, Professor M. Schwartz of Stanford University (SLAC) is developing parallel processing hardware which can be attached to a PDP-11 processor via the unibus. This hardware allows multiple operations to be performed at a high speed. One can add two 16-element vectors together as one parallel operation. As well, one can compare 256 data points with one common value and determine which of these 256 entries is greater than this common value. This special purpose processor is extremely inexpensive and compact. Yet it provides high speed execution of many parallel operators such as those found in the APL language.

- (d) Microprogramming techniques. Currently, several small computers permit the modification of their instruction set utilizing micro-programming techniques. For example, the Hewlett-Packard 2100 permits the extension of its instruction set by reloading alternate or modified versions of the basic micro-program. We visualize many applications in the area of data processing, communication control, and data compression which would gain substantially from new instructions specifically designed to improve the throughput of these operations. Previous studies of the implementation of microprogram techniques have indicated one to two orders of magnitude improvement in the handling of special problem areas such as communications control.
- (e) Interactive graphics hardware. The development of low-cost interactive graphic hardware is anticipated in the near term. Today, a display scope, vector generator and keyboard costs on the order of \$6,000. We would expect this cost level to be reduced to the \$1,000 to \$2,000 range in roughly three years. The primary reason for the price reduction is mass production coupled with lower component cost.
- (f) High speed remote links. High speed remote links will become available as conventional communication industry services are extended. Current networks such as APRANET and TYMNET are initial entries into this general utility field. The Bell Telephone System is currently testing a 500,000 bit per second line and will offer experimental lines at this rate. We do not contemplate early connection to such networks, but are watching their evolution closely as a future option.
- (g) Time multiplexed ring. The introduction of a commercial time multiplexed ring for digital communication started approximately three years ago in the form of the Collins C-System. This concept carries the potential of offering very high data rates at very low costs for inter-device communication. The data rates are sufficiently high that one could consider sharing memory among several satellite processors. It also provides a framework in which a number of satellite systems can be assigned selected portions of a common task. This concept may provide a redundant communications path among a large number of instruments and satellite machines throughout the Medical Center. A related development could well be a simple integrated circuit interface making it possible to inexpensively connect devices to the ring.

Perhaps, smaller subrings could be installed to individual laboratories. This ring would interconnect many experimental devices to one or several of the satellite computers. We believe that such a system would greatly reduce the problem inherent in interfacing large numbers of individual experiments to small satellite computers. The time division multiplexed ring may be compared to the PDP-11 unibus. The control features and inter-device communication paths are quite similar. The major difference is in the time division nature of the proposed ring which simplifies multiple

highspeed communication over a single path and also requires only one interconnecting wire rather than the number of wires in a PDP-11 unibus.

A time multiplexed ring on coaxial cable would have desirable characteristics for connecting satellite computers. This would provide a daisy chain in which very high data rates could be accepted. It would also provide a very reliable hardware connection. In addition to the possibility of replacing fixed head disks with core on a ring, one can hope to replace small expensive fixed head disks with a few large moving head disks.

- (h) Reliability. One general result of technological innovation will be substantial improvements in hardware reliability. This is a key point relating the use of satellite systems to the medical environment. Improved reliability can significantly impact the architecture of future medical systems.
- (i) Satellite software. As the problems of inter-device communications approach solution, we could expect to see satellite processing units designed to handle higher level languages efficiently. There are several examples of powerful high level language processors currently available from computer manufacturers. The Burroughs B 1500 provides specialized micro-computer support for COBOL, ALGOL, and systems programming applications. The SYMBOL machine developed by Fairchild utilized many cooperating processors to provide a high level language environment. These processors were dedicated to the tasks of syntax analysis, storage management, garbage collection, and input/output handling. Also, commercially available is the Hewlett-Packard System 3000. It too provides a machine organization optimized to support a higher level language processor. A processor for APL is under development at the University of California at Berkeley. This system utilizes special microprogram techniques on a Digital Scientific Corporation META 4 computer.* Each of the systems we have mentioned attempts to reduce overall software costs by incorporating special features and language oriented architecture within a satellite computer.

(b) Background and Rationale

- (1) Current environment for satellites in S. U. Medical Center. Stanford Medical Center today has approximately two dozen small stand-alone computing systems. They range in size from PDP-8's with 4K of core to a Sigma 3 with 32K of core. Some systems are being established to provide routine service. Examples of this include a Sigma 3 in the Clinical Laboratory, PDP-11's for Drug Interaction service in the Hospital Pharmacy, and an HP 2116 for physiological monitoring in the Catheterization Laboratory, and an HP 2100 being prepared for monitoring in the Cardiac Care Unit. Some of the research applications which currently use mini systems include PDP-11's for control of data acquisition and mass spectroscopy,

* "A Firmware APL Time-Sharing System", AFIPS-Spring 1971, R. Zaks, D. Stien-gart and J. Moore, U. C. Berkeley.

two PDP-12's for realtime data acquisition research in Cardiovascular Surgery and Psychophysiology, a number of PDP-8's and an HP 2116 for research in Psychiatry and Cardiology, and an HP 2100 for realtime data acquisition from nuclear cameras for the Division of Nuclear Medicine. A number of these systems will be connected to the ACME system during the summer of 1972 via a new small machine multiplexor interfaced to the 360/50 by a 16 bit PDA port on the 2701. A few have been connected to ACME in the past via the IBM 2701 and 270X. Research groups in the Medical Center have hired approximately 20 scientific programmers (exclusive of the ACME facility staff) to work primarily on new applications. In addition the ACME Computing facility has a number of programmers well versed in small machines. ACME is used as a consulting source by a number of the small machine users in the Medical School.

- (2) Initial approach in software. The small machine support development effort will initially be concerned with providing sophisticated software support for existing hardware in the Medical School. Specifically, it is our intention to provide a large number of new services for PDP-11 users and HP 2100 users. The new services will be software extensions to the systems currently provided by DEC and HP. For example, satellite machines will be given the ability to read and write files on the disk attached to the large host system, to communicate with other satellites in the network, and to prepare most of their software from a terminal connected to the host.
- (3) Incorporation of new hardware techniques. In addition to the task of adding software to better serve existing satellite computers, it is our intent to integrate the new hardware technologies listed above within the Stanford medical environment. One possibility would be the modification of the HP 2100 system using micro-programming techniques.

Historically, the manufacturers of small computing systems have been slow to provide software support for new hardware technologies. Although DEC has announced its PDP-11/45 with memory segmentation hardware, no operating system has been announced to make use of this new feature.

One reason for selecting the PDP-11 as the primary small system to be supported under this program is the expectation that the PDP-11 family will continue to be the primary focus of DEC support over the next several years in the minicomputer market and that it will have broad acceptance. This market acceptance in turn will lead to new technological innovations being made compatible with this particular line of hardware. Two examples of hardware support needed for the PDP-11 are a) the ability to execute programs in PDP-10 core from a PDP-11 and b) the ability to utilize PDP-11's as links in a network.

C. G. Bell of Carnegie-Mellon is currently developing a computer network utilizing a large number of interconnected PDP-11's which share one common extremely large primary store.* This system is being

* "C.mmp: The Multiminiprocessor Computer", C. G. Bell, W. Brody, W. Wulf, and A. Newell, Dept. of Computer Science, Carnegie-Mellon University.

constructed in order to provide an appropriate computer system for the realtime processing and recognition of human speech. Their proposed system appears extremely economical and well thought out. It provides a highly reliable environment in which a large number of small processors can intercommunicate and cooperate in the processing of a realtime analysis. The results of their research may have a significant impact in understanding the benefits or problems in interconnecting large numbers of small satellite processors.

Finally, the overall goal of intelligently managing large volumes of source data in a realtime environment, and processing large bursts of data with highly variable content and significance, will require coordinated and sophisticated use of small computers operating as satellites to the large host machine. No major computer manufacturer appears to place the solution to this problem high among the list of software support goals. For example, much of the realtime support being developed in the computing industry today is designed to support industrial manufacturing processes which do not in general have high burst data rates and sophisticated computational requirements.

- (4) Resolution of specific goals with collaborators. Early in the process of scheduling specific small machine support tasks, we will engage ourselves in extensive discussions with collaborators to determine which of their needs carry highest priorities and which of our proposed services should receive highest priority. A firm understanding of the user's environment is essential to the small machine support team. It is our intent to avoid the common pitfall of finding the solutions for which no known medical problem exists. This can only be done by close cooperation and collaboration with our prospective users. No large research tasks will be undertaken in the area of small machine support until a collaborator has been identified with a specific set of requirements to be met. This position does not prevent the development team from attempting to raise the horizons or expectations of the potential users.

(c) Methods and Procedures(1) Satellite system software support

(a) Assemblers. The advantage of providing assemblers on the large host machine is that satellite users are permitted to flexibly save, edit, and assemble source programs for new applications. In addition, they have the advantage of better error diagnostics in far less time than is the case on their dedicated systems. At present the ACME system supports assemblers for the PDP-11, PDP-8 and the IBM 1800. (This will be extended to include the HP 2100.) The existing assemblers will be updated from time to time to take advantage of the extensions to the manufacturer-supplied operating systems as described below.

(b) High level language support.

(i) Machine dependent. We can simplify the code generation process while maintaining highly efficient execution on the satellite processors by introducing new languages modelled after PL/360 or BLISS. The applications programmer will be able to obtain optimal machine code without the need to be concerned with all details of architecture of the hardware upon which he will be running. The programmer is not prevented from inserting machine language where he feels he can improve upon the facilities provided. Over a five-year period, we would expect to spend roughly two man-years on this effort. Carnegie Mellon has produced a BLISS-11 and CERN Laboratory in Switzerland is producing a PL/11. We will select the best support from among several such development efforts and make them available at the Stanford Medical Center. We will not concern ourselves with the provision of standard language support in small machines as this is already being undertaken by manufacturers and others.

(ii) Machine transparent. The second class of non-machine-oriented language support deals with non-standard languages such as APL, SNOBOL, and LISP. Providing a subset of APL is highly appropriate for satellite machines. We intend to define new APL-like primitives and couple them with a subset of existing APL primitives to provide a high level language support for realtime medical computing. We want to make available an interactive language with user-oriented runtime diagnostics, simple language structure, and the ability to modify programs without complete recompilation or assembly, coupled with graphic output. Some examples of such primitives are:

a. Smoothing primitive. This primitive would employ a number of arguments such as window size, resolution, convolution function, or other user-supplied arguments. The second parameter would be the new data. The primitive would return smoothed data as its value.

- b. Peak Extraction Primitive. This would return the coordinates of all peaks found above a threshold provided as a parameter by the user. If the data provided is a vector, the primitive would return the set of index values at which the function attains its peaks. If the data are paired coordinates (the first being the data point and the second being, for instance, time) then the primitive would return the peaks and times, possibly interpolated between two time values.
- c. Curve Fitting Primitive. The parameters associated with this primitive would be data and name of a system or user function which would provide the shape of the curve to be fitted. The primitive would search for the parameters to the function which provide the best fit among the user supplied data points. The user function will automatically handle fitting of overlapping peaks as a natural outgrowth of the nature of the APL language.

Some of the current APL operators which would be very useful in a realtime environment include matrix manipulation, sorting, bit manipulation, and conversion of raw BCD data to binary by using the encode operator.

- (c) Realtime applications module library. Some of the items expected to be provided are the following:
 - (i) Data collection routines. These would include automatic sampling of selected data channels at user specified breaks, interrupt driven data collection, and automatic scheduling of resources on the inter-computer communication network.
 - (ii) Data compression routines such as the Aztec procedure developed by Dr. Jerry Cox for compression of EEG data at Washington University in St. Louis.
 - (iii) Input/Output buffering routines. These routines would automate core management since, due to the great variability of data rates, a sophisticated scheme is necessary. Extensive effort will be needed in this area because manufacturer-supplied operating systems have not addressed themselves to the high data rate, burst mode, realtime data collection problem. The intelligent management of such data streams will require innovative techniques in core management.
 - (iv) Peripheral handlers. We wish to provide a standardized approach to unique peripherals similar to that currently provided by the manufacturers for standard peripherals. The newly developed peripheral handlers (for instruments such as mass spectrometers and gas chromatographs) can then utilize a standard interface to the operating system.

- (v) Communications protocol with host computer. User programs would interface to the communications network via routines present in the applications library. The detailed operation of the communications system need be no more apparent to the user than the details of a disk or tape controller would be.
- (vi) Debugging packages. Along with all the new facilities outlined above, we must provide the user with a complete set of debugging packages for realtime control systems. This function is essential in any realtime environment, where any one of several links in the total chain may prevent satisfactory operation. The ACME facility has consistently found it difficult to identify just which link in the chain is defective (despite the best efforts of good engineers and programming talent!).
- (d) Extensions to manufacturers supplied operating systems. Handling realtime functions and using multiple processor configurations are two examples of activities contemplated here which are not covered by current operating systems but are needed in our labs. Our computing objectives cannot be met without extensions to what exists. For example, DEC permits several PDP-11's to be connected via Unibus links, but no multiprocessor software is provided for this (marketed) configuration. The particular problem areas where extensions are required have already been described above.

(2) Computer to computer communications.

(a) Data transfers.

- (i) To share peripheral equipment. Software support should exist to make possible the sharing of peripheral devices on a host computer with all satellite processors. For example, high speed printers on the host should be accessible to data collections or data analysis programs on satellite systems. Similarly, results of data analysis on host machines should be available at printers connected to satellite systems. It is hoped that our commitment to software support for shared peripheral equipment will reduce the overall number of peripherals required in the Medical Center.
- (ii) To move data to other systems. Data transfer software support will facilitate transfer of data collected on mini-computers to the host system and vice versa. In some cases, an intermediate system will serve as a store and forward device. In this way users will not be dependent upon the availability of the large host machine for long term data collection problems. A specific example of data transfer among satellite systems is the need for sharing of information between Clinical Laboratory and Pharmacy systems. The interpretation of a certain laboratory result could be influenced by the knowledge that a given drug had been administered prior to collecting the laboratory specimen.
- (iii) Loading of Object Modules. Software support will be developed so that object modules can be loaded to satellite machines from the host machine. In this manner, satellite machines avoid the necessity of loading object modules from paper tape or other more expensive peripherals.

(b) Controlled interactions.

- (i) Start satellite machine (IPL) remotely. Code will be developed to permit the startup of a satellite by a program in the host machine. The satellite in turn could be programmed to start up various laboratory instruments. The benefits of such a procedure are to permit control of experiments from one terminal and to enforce a common chain of events in the startup of each day's operation in the laboratory.
- (ii) Detect failure of other systems. The host machine can be programmed to detect failure of satellite systems and satellite systems in turn can be instructed to detect failure of instrumentation. In the case of communications systems, failure detection is essential so that control of the communication system can be passed to an alternate source. The same will be true for some physiological monitoring systems.

- (iii) Network management. Comtec, Honeywell, CDC, and Tymshare have devoted considerable effort to programming intelligent machines to handle network management tasks. It is not our intention to duplicate this effort. However, we may wish to supplement the existing network management programs by adding realtime support functions. One possible step in networking will entail the installation of a long-distance attachment for the ARPANET IMP. An IMP exists on the Stanford campus at the Artificial Intelligence Laboratory directed by Dr. John McCarthy. The long-distance attachment for an IMP is being designed at U. C. Santa Barbara. Assuming that the design will work, it may provide a relatively inexpensive connection to the ARPA Network. The Hardware costs exclusive of "lease" line will be approximately \$15,000. At this time, we have elected to study the networking opportunities without further commitment.

(3) Data manipulation.

(a) Data collection

- (i) Feedback to local controllers of realtime data flow. Intelligent management of large volumes of data implies an interaction between the initial data received, the long term goal, and feedback to the experimental apparatus. The closed loop situation is one in which special benefits should accrue as a result of the small machine support effort. Development of heuristic processors in models should permit increasing amounts of feedback to local controllers for instrument adjustment, calibration, and data management.

New predictive techniques could aid in image processing. In classical digital image processing, an image is first digitized to be in a machine readable form. This requires 10^5 to 10^7 bits, depending on the requirements for resolution and grey scale. At this point, data compression methods are often applied to reduce storage and computational requirements. This may result in compression factors of two to ten. These actions require additional computation and only serve to make the data manageable, rather than to extract information from it.

Another approach is to record only those elements of an image which differ from the previous image. This reduces the amount of data but still does not recognize the fact that for a large class of images (e.g., a beating heart) the position of the elements in a new image can be predicted with a high degree of accuracy, knowing the time between images, scale factor, and direction over the past few images. Since this is true, a significantly reduced number of data points need be taken to refine both the prediction and the model used to make the prediction. Gross deviation from a well established model might be used to indicate a pathological condition.

It is felt that this approach to image processing helps to reduce the gross data handling requirements to a manageable level. It concentrates required complex image analysis algorithms on that subset of the raw data worthy of attention. By using such methods, we feel that significant results can be drawn for a class of problems of biological interest.

- (ii) Multiple paths to large files. We plan to provide multiple paths to large files from satellite computing systems. In this way, we can guarantee access to the large files for a 24 hour day, seven days a week period. Only by providing assurance on this point can satellite machine owners be convinced to assume the risk of minimizing their local dedicated configurations and depending upon shared use of larger resources.
 - (iii) Store and forward system. One satellite in the system may be dedicated to realtime communications management and interim storage of realtime data. The primary rationale for having such a facility is a) to guarantee better response time, b) to provide larger storage capacity than would be feasible for each of the experimenters connected to the system, and c) to assure continuous availability either from the host or from the store and forward processor.
- (b) Data analysis.
- (i) Multiple processor allocation. Optimal use of the resources available in a large system implies that repetitive tasks should be handed to less expensive processors dedicated to a more limited range of tasks. It is our intent to demonstrate that selected subroutines which are frequently used in a realtime environment can be passed to a satellite processor for execution, thus freeing the host machine for other tasks.
 - (ii) Software support for special hardware. Special purpose devices such as Fast Fourier Transform hardware, matrix multipliers, and graphics aids can be supported for satellite systems. One can make a specialized device available to more than one user's laboratory while at the same time limiting the software investment for its use.
 - (iii) Interactive graphics. Graphic interactions with realtime data streams is essential in the rapid interpretation by man of "what's going on". The objective of this effort will be to provide a smooth human interaction with graphics devices in both directions (man-to-machine and machine-to-man). The anticipated decrease in cost of graphics stations is likely to increase significantly the demand for this service. Common generalized software for support of such instruments can provide an important savings in programming manpower.

ACME currently supports a variety of graphics devices in a flexible manner. All graphics devices are treated as real-time output devices. Associated with the description of output devices on a realtime line are the programs that convert a graphic description into the detailed control sequences required for device operation. The user needs only to change the destination number parameter in his graphic output calls and the identical user program will drive any of the graphics devices available at ACME.

The devices currently supported are:

- a) The ACME TV display (high precision, refreshed by an auxiliary core memory).
- b) Tektronix 611 storage tube displays.
- c) Calcomp and Houston incremental plotters.
- d) General Purpose Graphics Terminals.

In addition to supporting other display types as they become available, ACME is planning to extend support for graphics activities in the following ways:

- a) Development of interactive generation of display programs. The abstraction of a planned visual image into display-driving statements is not always obvious or easy for medical researchers. Computer languages such as PL/ACME, while providing all the required capabilities, do not express the two dimensional nature of graphs very clearly. A question and answer communication between computer and user is being prototyped which will generate the required instructions using decision trees which systematically reduce the alternatives of choice. The resulting protocols can be saved and subsequently modified as desired.
- b) Current display support has been mainly oriented toward graphics. The General Purpose Graphics Terminal has the capability to handle text. This facility has been made available now to General Purpose Graphics Terminals users, but it should be made equally available on the other devices supported in graphics mode, which now have only limited text capability.

- (4) Satellite performance measurement and evaluation. Very little has been done in the area of performance measurements for small computers. We feel that there is much to be gained from activity in this area. Both hardware and software monitors will be investigated.

It is known that most programs spend a large amount of time executing a small percentage of their total code. The classic approach to increasing a program's efficiency is to locate the most used code and

recode it in the most efficient manner available. If a program is to be written in a higher level language, the ability to "tune" with performance measurement tools may make code generated by a sub-optimal compiler (which many small-machine compilers are) acceptable.

Such monitors are also valuable debugging aides. The sorting out of a sequence of randomly occurring realtime events can be an almost impossible task without some monitoring ability. Excess time spent in error recovery can be located and corrected. Inadvertent loops may be detected by monitoring.

It is felt that this type of support for the satellite machine falls into that class of activity which needs very much to be done, but which the individual mini user cannot afford to do.

(d) Significance of satellite machine support.

- (1) Remove the limitations in users' laboratories: Current satellite machine users continuously encounter physical limitations in their hardware, such as lack of core, lack of disk space, lack of registers, inadequate cycles, lack of generalized data acquisition, reduction and analysis subroutines, and problems associated with interfacing hardware to experimental apparatus. The proposed research program and satellite machine support will help to overcome most of these limitations.
- (2) Reduce effort required of experimentalists: The application of standard interface hardware to the extent possible will reduce the amount of active involvement of the experimentalist in this problem. Furthermore, availability of tried and tested software for handling many of the problems frequently encountered in data acquisition situations will allow him to select from a library those elements which he needs. Additionally, the local satellite machine group will be aware of routines developed elsewhere in the country through DECUS and other user groups. As the relative investment in software becomes far greater than hardware, the impact on the user of extensive satellite machine support will be great.
- (3) Effective easy access to shared data: Computers can truly be used as a means of sharing knowledge when the large volumes of machine readable data can be shared with trivial effort on the part of the user and at acceptable costs. Sharing of data among Clinical Pharmacology, Hospital Pharmacy, Infectious Disease Laboratory, other clinical laboratories, and physicians is currently being requested but has not been provided other than through manual transport of magnetic tape. In the research area it is felt that new algorithms, models, new graphics techniques, and other developments will be shared throughout the community more quickly with improved communications and file access systems.
- (4) Integrate benefits of small and large systems: Both the small and the large systems offer unique benefits. The objective is to realize the synergism believed possible through the marriage of small and large systems.

(2) Extended Realtime Computing(a) Problem Statement and Objectives

State-of-the-Art - Applications of computers have developed over the past ten years which involve the interaction of computer systems with various aspects of medical research. Computers are used in a broad spectrum of ways in support of the acquisition of raw instrument data, the reduction and standardization of data quality, the interpretation of experimental information and building of models, and finally the design of new experiments to test the models. Historically computers have performed simple supporting tasks such as prescribed data logging procedures and reduction computations with human investigators carrying out the higher level processes of adaptive system control as well as theory building and testing. In most cases errors arising from simplistic computer processing of the data are detected and corrected by human intervention. Typically the computer does not have access to a model of what it is doing to evaluate its success or failure. In applications involving relatively small amounts of data which can be collected simply and analyzed without severe time constraints, this division of labor is satisfactory and economical. Indeed this separation is to some extent necessary because computer programs are currently incapable of many of the complex reasoning and creative problem solving processes necessary for medical research.

Limitations and Future Needs - There is an increasing number of situations, however, for which this type of open loop or loosely closed loop solution is infeasible or unacceptable. The requirements for more automated loop closure may grow from a variety of circumstances such as:

- (1) Human boredom with making large quantities of precise, detailed measurements.
- (2) Experiment time constants allowing the collection of only a subset of possible types of available information requiring judicious on-line selection.
- (3) Economy of extracting small sets of significant information from large quantities of raw data.
- (4) Adaptive experiment optimization and control requirements too complex or rapid for human response.
- (5) Large and complex information bases or models difficult for human manipulation and analysis.

A variety of examples can be listed of applications facing these problems today and most certainly increasing demands on computing capabilities of this kind will arise in the future. Such fields include image processing and perception (microscopy, radiology, ultrasonics, etc.);

stimulus/response experimentation (neurophysiology, anesthesiology, etc.); analytical instrumentation (gas chromatography/mass spectroscopy, x-ray diffractometry, etc.); and interactive system and data modelling.

Intelligence Requirements - Characteristic of these problem areas is a requirement for more and more "intelligent" handling of information within appropriate time constraints. The term "intelligent" is used to imply autonomous and adaptive performance of necessary operations. In becoming autonomous the computer must take advantage of the dynamic characteristics of the input data and extracted information as well as previous or evolving problem solutions to economically produce accurate and reliable results. In the simplest form, the computer has a model of its environment relative to the task it is performing and uses derived measures of success or failure to optimize performance. In the longer term more sophisticated reasoning and inductive processes are applied to an information base.

Difficult conceptual and implementation problems exist in designing intelligent information handling programs for computers. Beginnings have been made in developing such capabilities in a few applications. Much more work is necessary to improve program capabilities and reliability as well as to explore wider application areas in medicine.

Computing Requirements - As the sophistication of computer processing of information increases, so does the requirement for intensive usage of large computing resources. These requirements are typically applied over relatively short or sporadic periods of time. In effect each such application requires a large dedicated capability for high demand on-line work and a less responsive time-sharing or batch support for off-line or developmental work. Thus the extended realtime user is faced with the dilemma of needing to control for short periods of time a machine of a capacity that does his job but which he cannot afford or justify having totally dedicated to his work. Methods for providing this type of service economically must be developed.

Research Objectives - The overall objectives of this portion of the proposal are to develop a computer resource to investigate a range of problems associated with extended realtime computer applications in medicine including:

- (1) The utilization of intelligent methods of information handling to increase system effectiveness and reliability and to allow the solution of problems which are unmanageable by brute force techniques.
- (2) Methods for organizing and delivering large capacity computing power to extended realtime users within required time constraints and within the context of complementary time-sharing and batch machine utilization.

The short term objectives will be to fashion solutions to specific problems among a small set of collaborators. In the longer term more general methods for dealing with larger communities of users may become apparent. Experience has shown that progress with automated computer systems is difficult with success depending on the careful selection of problems as well as techniques. It is recognized that considerable differences exist and indeed must exist in hardware and software needed to solve specific problems. It is not our goal to attempt to force the solution to all problems within a rigid framework but rather to exploit the aspects of commonality between applications while explicitly allowing for necessary differences.

(b) Background

Realtime Experience - Many laboratories are working on developing on-line experiment support. At Stanford the ACME computing resource has had as one of its objectives the development of on-line, realtime instrument support capabilities in a time-shared environment (ref. 4,5). This system has had success in servicing laboratory instrumentation primarily as a store and ~~forward~~ data logging service followed by near realtime reduction processing. Experience has been gained in a variety of on-line applications including the measurement of heart function parameters (refs 6, 9-12), respiratory function parameters and interactions (refs 1,14), electroencephalogram correlations (refs 7-8), urinary function parameters (refs 2-3), and the design and application of gas chromatograph/mass spectrometer data systems (ref. 14).

Automated Systems - A number of application problems have arisen in the course of this work where significant autonomy in the computer handling of information is required. Specific examples include the quantitative analysis of cineradiograms (cardiology and urology), electroencephalograms, and mass spectograms (see the succeeding collaborative experiment descriptions). Without computers the quantitative utilization of these sources of information would be infeasible. Our experience in the design of automated mass spectrogram interpretation systems (see Section G.2.b.3. for Dendral references), and image processing systems (refs. 21-22, 31-32) indicate two types of problems will be important: First the design of algorithms to accomplish the required information extraction and second the design of algorithms to assess system performance to optimize information quality. Most such systems are complex and display highly abstracted versions of the source information as results. Unless models are available to check the quality of intermediate processing stages, subtle losses or distortions of the information may result leading to misinterpretations. Both types of problems are difficult, application specific, and require intensive computing resources.

Closed loop problems are under consideration in a variety of applications either with largely dedicated computer systems or with shared machines where loop closure time constants are flexible. Examples include image processing and pattern recognition systems, analytical instrumentation systems, robot systems, spacecraft operations, and military systems.

The DENDRAL programs at Stanford are able to infer the structures of complex biological molecules from mass spectra without human aid. Further work is under way to automate the computer extension of the domain of mass spectral problems it can solve. Highly autonomous robot systems such as SRI's SHAKEY (ref. 23) and the Stanford Artificial Intelligence Hand-Eye Project (ref. 24) are able to completely solve "simple" problems based on initial goal selection within geometrically structured environments (such as stacks of cubes). Pattern recognition systems dealing with more amorphous environments such as bubble chamber tracks (ref. 25), karyograms (refs. 26-28), or radiograms (refs. 29-30) must either carefully select samples for processing or rely on manual intervention. Much work remains to be done in automating computer applications to medical problems. We will draw upon related research as appropriate in our work.

Impact of Small Machines - Closed loop experiments on ACME have been limited to relatively low rate interactions or non-real time closure because of capacity constraints and usage priority conflicts. These difficulties with central machine support have led to an increased use of "mini" computers dedicated to each particular laboratory environment. For a variety of economic and technical reasons this experience is common as witness the tremendous market which has built up around small computers over the past few years. The trend of servicing laboratory instrumentation with local minicomputers is unmistakable and numerous self-contained instruments with supporting data systems are available commercially.

Small Machine Limitations - Small dedicated computer systems provide highly flexible data logging devices and perform straightforward data reduction tasks as well. A machine restricted in memory, peripherals, and instruction flexibility, however, must exercise many short-cuts to achieve results. These short-cuts restrict system adaptability and limit the domain of data-interpretation algorithms which can be employed.

Such limitations are offset by incrementally expanding "small" systems by adding special features such as floating point instructions or relatively expensive memory, disk, and other peripherals. The result is an increasingly expensive piece of general purpose equipment dedicated to a task, and inefficiently used.

Our experience has shown that such data systems, built around small machines, are adequate if limited realtime demands exist. The presence of a human being in the loop to make control decisions and adaptive adjustments is essential in many situations. More demanding applications, however, require the examination of more parallel system hardware and software configurations designed to balance the economy and capacity of central and distributed machines.